


discontinue



Massachusetts Institute of Technology
Artificial Intelligence Laboratory

AI Memo 406

Log Memo 40

Replaced by
Laurie's Paper

Overlays: a Theory of Modelling for Computer Aided Instruction¹

Brian Carr and Ira P. Goldstein

February 1977

Overlay modelling is a technique for describing a student's problem solving skills in terms of a modular program designed to be an expert for the given domain. The model is an overlay on the expert program in that it consists of a set of hypotheses regarding the student's familiarity with the skills employed by the expert. The modelling is performed by a set of P rules that are triggered by different sources of evidence, and whose effect is to modify these hypotheses. A P critic monitors these rules to detect discontinuities and inconsistencies in their predictions.

A first implementation of overlay modelling exists as a component of WUSOR-II, a CAI program based on artificial intelligence techniques. WUSOR-II coaches a student in the logical and probability skills required to play the computer game WUMPUS. Preliminary evidence indicates that overlay modelling significantly improves the appropriateness of the tutoring program's explanations.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. It was supported in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643, and in part by the Division for Study and Research in Education, Massachusetts Institute of Technology.

¹ This paper has been submitted to the Fifth International Joint Conference on Artificial Intelligence.

1. The MIT COACH Project

A traditional argument for computer aided instruction (CAI) has been that it is an economic means for providing individualized instruction. The rapidly falling costs of hardware make the economics of CAI progressively more appealing. But, the extent to which existing CAI has provided personalized instruction has been limited. This paper develops a procedural theory of modelling that can be incorporated into CAI programs to address this limitation.

This theory has been developed as part of the COACH Project at MIT, whose concern is the development of AI-based CAI programs for tutoring the skills required for successfully playing various computer games. The computer serves as an assistant to a learner who is in the process of acquiring the skills necessary to play the game well. Fig. 1 shows a generalized block diagram for these programs, with the modules given anthropomorphic names to indicate their function. To distinguish them from their human counterparts, references to the modules will be capitalized.

Good coaching is critically dependent on a detailed model of the learner in that the model guides the coach in generating concise and appropriate explanations. This paper discusses the theory of overlay modelling embodied in the Psychologist module, the component of the Coach responsible for maintaining such models of a player's current skills (the K model) and learning preferences (the L model). These models are used by the Tutor module to prune complex explanations generated by the Expert. Just as with a human speaker, the Coach abbreviates its statements by eliminating those facts that are already known by the listener and those facts which are too complex

A broad treatment of the potential role of computer coaches in education and the issues raised by their design is given in [Goldstein 1977]. Detailed discussions of preliminary implementations and experimental results are provided in [Stansfield, Carr and Goldstein 1976] and [Carr 1977]. Seminal work on AI-based CAI is also described in [Brown et al. 1975; Brown 1976; Collins & Grignetti 1975]. In particular, overlay modelling is an extension of the issue-oriented approach to student modelling developed by Burton and Brown [1975].

Overlay modelling is a technique for recognizing the constituent skills being

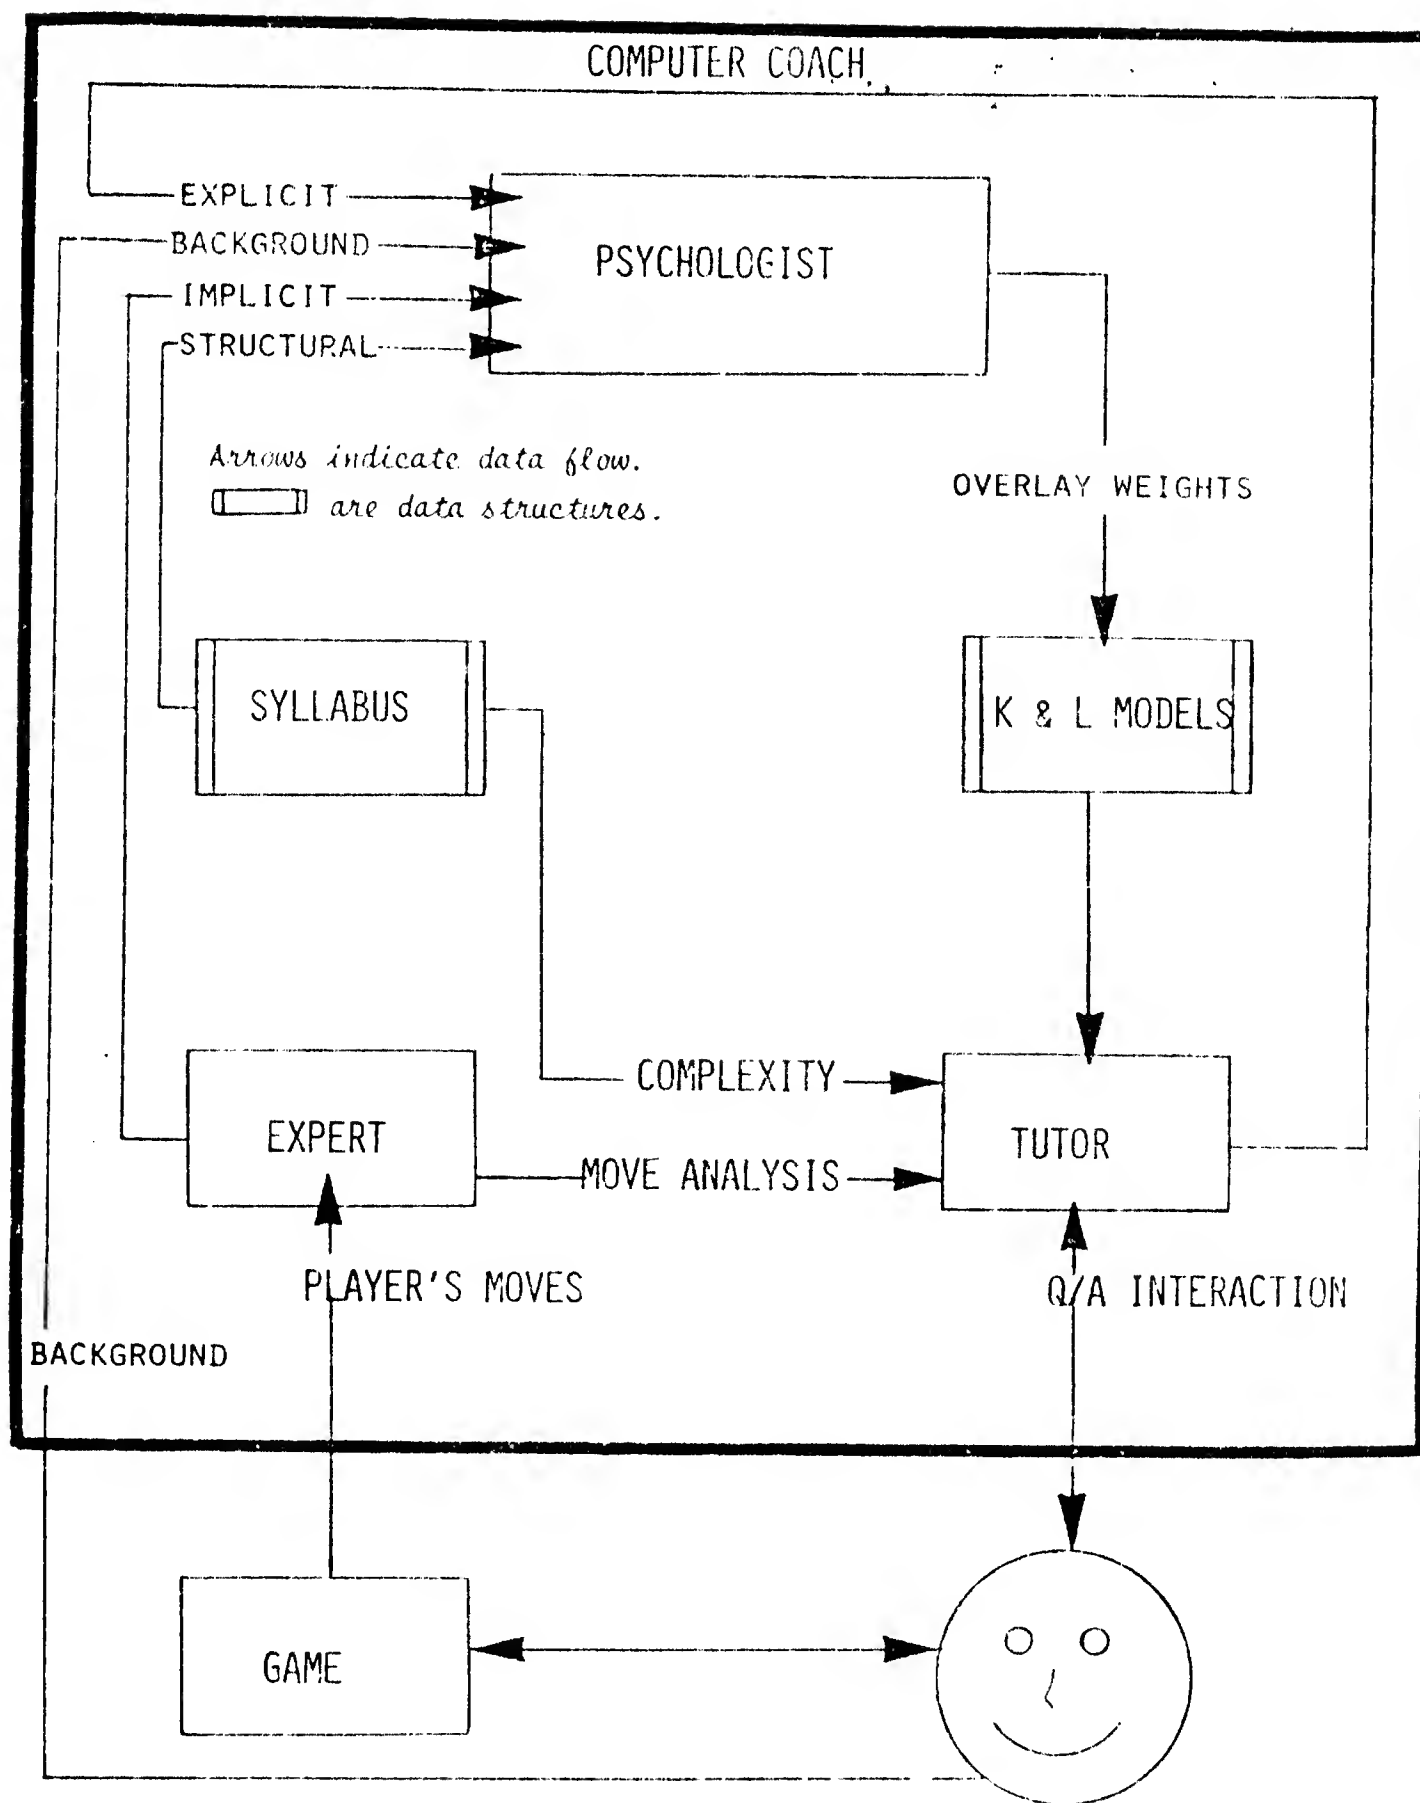


FIG. 1 -- BLOCK DIAGRAM OF A COMPUTER COACH

exercised by an individual in performing a problem solving task. The kernel idea is to design a modular Expert program for the task, and to explain differences between the behavior of the Expert and the subject in terms of the lack, on the player's part, of some of the Expert's skills. Thus, a model of the player is a set of hypotheses, each of which records the system's confidence that the player possesses a given skill. Such models are called overlays to reflect that fact that the model of the individual is basically a perturbation on the Expert's structure.

Overlays in terms of subsets of the Expert's skills is a simplification of the modelling problem in that it does not address situations in which the student has an incorrect skill or an alternative skill. A discussion of this limitation is given in section 5.

Modelling a learner is difficult. However, preliminary evidence with WUSOR-II indicates that, at least for the restricted environment of a game and for the limited purpose of guiding a tutor, adequate modelling can be obtained from: a rule system that accesses multiple sources of evidence, and a critic that detects inconsistencies and discontinuities in the player's behavior. Fig. 2 is a block diagram of the internal structure of the Psychologist.

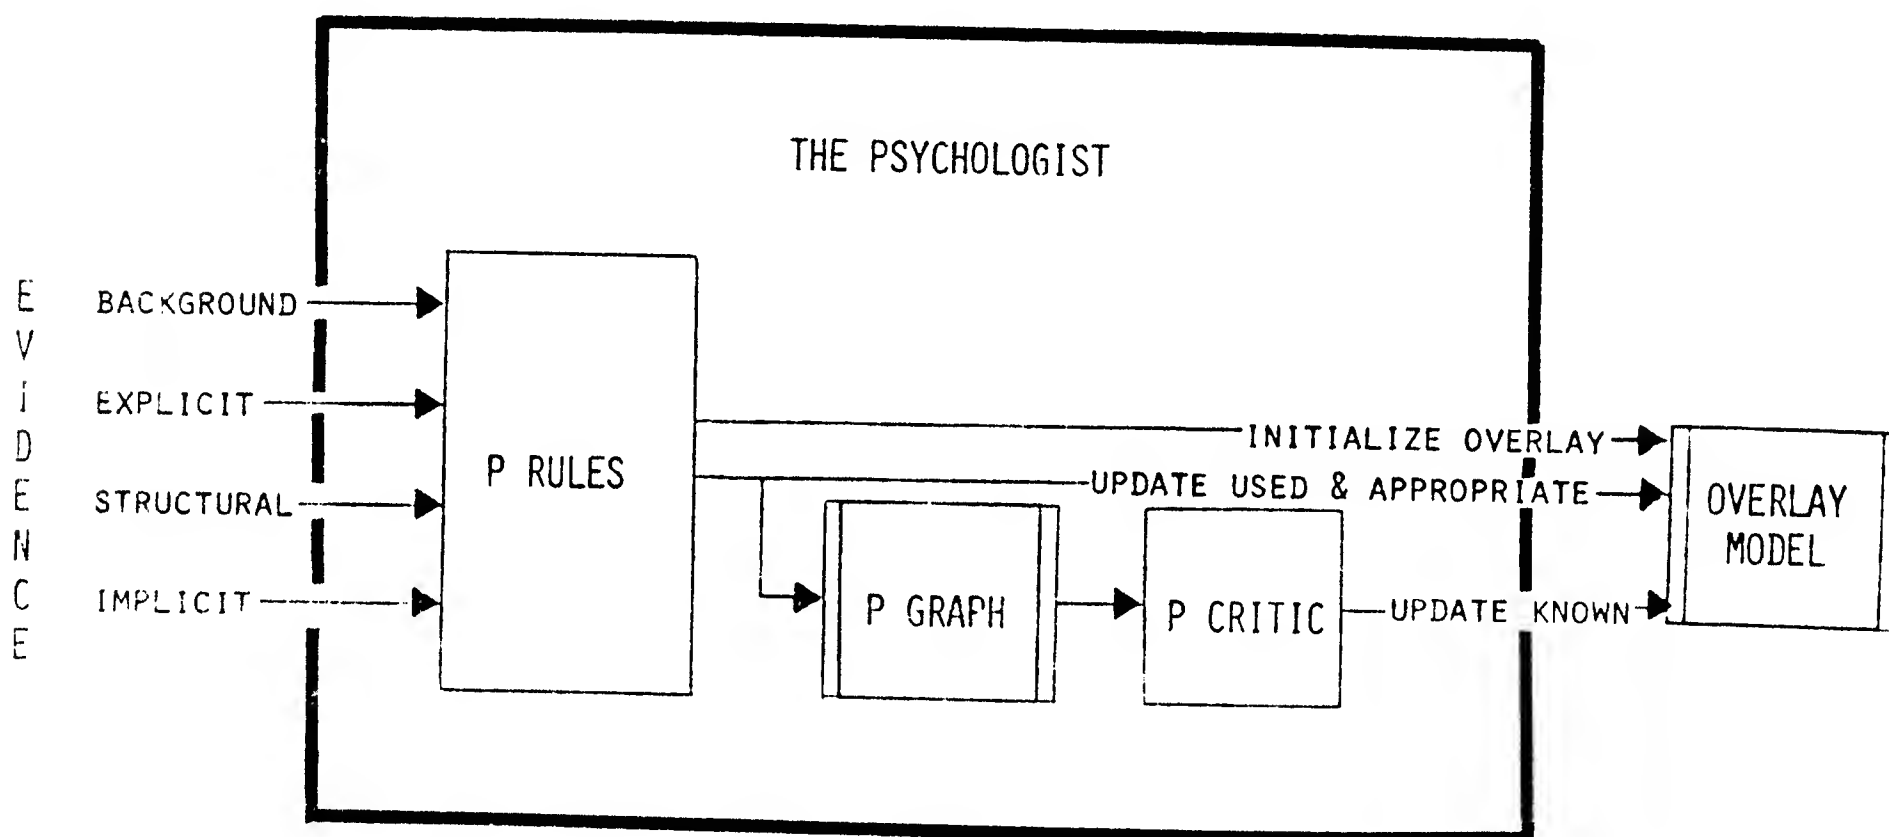


FIG. 2 - INTERNAL STRUCTURE OF THE PSYCHOLOGIST


Section 2 describes the Wumpus game, the experimental domain of the WUSOR-II coach. The theory of overlay modelling is developed next (sections 3-4), followed by a discussion of its limitations and extensions (sections 5-7), and concluding with our experimental program and preliminary results (sections 8). Related literature is surveyed in section 9.

2. Wumpus, an Intellectual Game

The Wumpus game was invented by Gregory Yob [1975] and exercises basic knowledge of logic, probability, decision analysis and geometry. Players ranging from children to adults find it enjoyable. The game is a modern day version of Theseus and the Minotaur. The player is initially placed somewhere in a randomly connected warren of caves and told the neighbors of his current location. His goal is to locate the horrid Wumpus and slay it with an arrow. Each move to a neighboring cave yields information regarding that cave's neighbors. The difficulty in choosing a move arises from the existence of dangers in the warren -- bats, pits and the Wumpus itself. If the player moves into the Wumpus' lair, he is eaten. If he walks into a pit, he falls to his death. Bats pick the player up and randomly drop him elsewhere in the warren.

But the player can minimize risk and locate the Wumpus by making the proper logistic and probabilistic inferences from warnings he is given. These warnings are provided whenever the player is in the vicinity of a danger. The Wumpus can be smelled within one or two caves. The squeak of bats can be heard one cave away and the breeze of a pit felt one cave away. The game is won by shooting an arrow into the Wumpus's lair. If the player exhausts his set of five arrows without hitting the creature, the game is lost. Fig. 3 illustrates a typical intermediate state a player might reach.

Skilled play exercises knowledge of logic, probability, decision theory and geometry. The WUSOR-II Expert uses a rule-based representation of this knowledge, consisting of approximately 20 rules, to infer the risk of visiting new caves. However, for expository purposes, a simplified rule set consisting of five reasoning skills is sufficient to illustrate overlay modelling.



- L1:** (positive evidence rule) A warning in a cave implies that a danger exists in a neighbor.
- L2:** (negative evidence rule) The absence of a warning implies that no danger exists in any neighbors.
- L3:** (elimination rule) If a cave has a warning and all but one of its neighbors are known to be safe, then the danger is in the remaining neighbor.
- P1:** (equal likelihood rule) In the absence of other knowledge, all of the neighbors of a cave with a warning are equally likely to contain a danger.

P2: (double evidence rule) Multiple warnings increase the likelihood that a given cave contains a danger.


In terms of these skills, an overlay model for a player who has mastered the simple logical rules (L1,L2), is in the process of acquiring L3, and has not yet learned P2 is:

| RULES | APPROPRIATE | USED | FREQUENCY | KNOWN |
|-------|-------------|------|-----------|-------|
| L1 | 5 | 5 | 100% | T |
| L2 | 4 | 3 | 75% | T |
| L3 | 4 | 2 | 50% | ? |
| P1 | 5 | 5 | 100% | T |
| P2 | 4 | 1 | 25% | NIL |

Overlay Model 1

The frequencies are determined by estimates made by the P rules of the number of times a skill has been USED in proportion to the number of times it has been APPROPRIATE. The KNOWN variable is set to T, ? or NIL by the P critic.

The WUSOR-II Coach [Carr 77] maintains models of this kind for guiding its explanations to the student. For example, consider fig. 4: Suppose the player moves



to cave 14, the worst possible move. Given overlay model 1, WUSOR-II would generate the following tutorial advice:

Ira, it isn't necessary to take such large risks with pits. One of caves 2 and 14 contains a pit. Likewise one of caves 0 and 14 contains a pit. This is multiple evidence of a pit in cave 14 which makes it probable that cave 14 contains a pit. It is less likely that cave 0 contains a pit. Hence, Ira, we might want to explore cave 0 instead.

Without the overlay model, the explanation would be longer and more complex as shown below. The WUSOR-II Tutor has pruned the underlined text from the Expert's complete analysis by noting that the student is already familiar with the positive and negative evidence rules.

Ira, it isn't necessary to take such large risks with pits.

Cave 4 must be next to a pit because we felt a draft there. Hence, one of caves 15, 2 and 14 contains a pit, but we have safely visited cave 15 This means that one of caves 2 and 14 contains a pit.

Likewise cave 15 must be next to a pit because we felt a draft there. Hence, one of caves 0, 4 and 14 contains a pit, but we have safely visited cave 4. This means that one of caves 0 and 14 contains a pit.

This is multiple evidence of a pit in cave 14 which makes it probable that cave 14 contains a pit. It is less likely that cave 0 contains a pit. Hence, Ira, we might want to explore cave 0 instead.

Thus, the overlay model has allowed the tutor to focus on explaining the double evidence heuristic to the player.

3. The P Rules

No single source of evidence is a certain indicator of an individual's knowledge. Hence, the Psychologist is provided with four sources of evidence -- (1) implicit (the student's behavior in playing the game), (2) structural (the intrinsic complexity relations between skills of the Expert), (3) explicit (the dialog between tutor and player), and (4) background (estimates of how average players of varying backgrounds can be expected to perform).

In this section, we define the P rules, a set of procedures which modify the overlay model when triggered by these various kinds of evidence. Section 4 describes the P Critic whose function is to set the KNOWN variable on the basis of the *history* of changes to USED and APPROPRIATE. In these sections, our example is the creation and maintenance of the K model, an overlay on the Expert. [Goldstein 77] describes the application of overlay techniques to the creation and maintenance of the L model, an overlay on the Tutor.

Implicit Evidence: The student's play yields implicit evidence regarding his mastery of various skills. The Expert evaluates the merits of the player's move relative to the available alternatives. The assumption is that the player has learned those skills involved in choosing his particular move and rejecting its inferiors, and has yet to learn those skills needed to recognize superior moves.

The implicit evidence rules utilize the Expert's analysis as follows:

P-I1: If skill S is involved in an overlooked superior move and not in the current move, then increase APPROPRIATE by $C(S)$ and recompute the frequency.

P-I2: If skill S is involved in the current move and not a rejected inferior, then increase USED and APPROPRIATE by $C(S)$ and recompute the frequency.

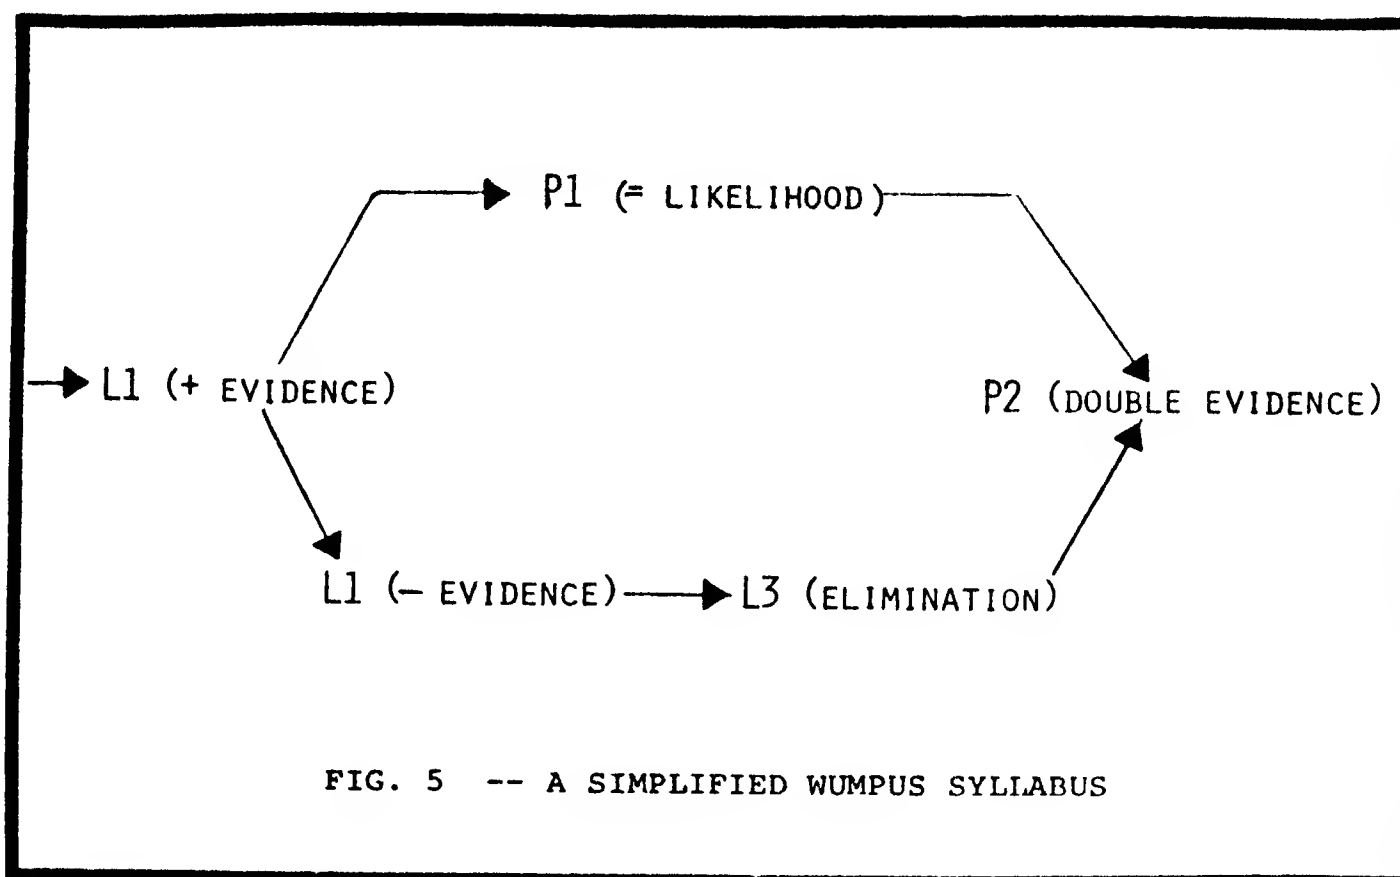
where $C(S)$ is a complexity factor ranging between 0 and 1 that decreases as the skill becomes more complex relative to the student's current knowledge state. $C(S)$ is defined in the next section

For example, in situation 1 the Expert reports to the Psychologist that caves 0 and 2 are better than 14 on the basis of double evidence (P2). If the player chooses 14, the Expert's analysis triggers P-I1 which increments APPROPRIATE but not USED. (The FREQUENCY of P2 therefore drops.) On the other hand, if the player chooses cave 0 or 2, then P-I2 would be triggered and both USED and APPROPRIATE for P2 would increase.

The Expert also reports to the Psychologist that cave 0 is better than cave 2

because of the known bat in the latter. Hence if the player chooses 2, P-II is triggered and the frequency of use of L3 drops, while choosing 0 has the opposite effect.

Structural Evidence: Clues to the student's knowledge arise from an analysis of the intrinsic structure of the skills to be conveyed. This analysis of the Expert's skills is stored as the Syllabus, a network linking the skills in terms of their complexity and dependencies. Fig. 5 is a simplified Wumpus syllabus for the five reasoning skills introduced earlier.



Structural knowledge suggests that given a student familiar with a certain region of the syllabus (as indicated by the K model), it is more likely that a new skill being acquired is at the frontier of this region rather than deep into unknown territory. WUSOR-II implements this heuristic in a conservative fashion: $C(S)$ is set to zero for every skill more than one away from a known skill. WUSOR-II thus ignores the possible employment of skills not at the frontier.

We currently believe that this is too conservative. It assumes that skills can only be learned in the order in which they appear in the syllabus. Such an assumption is too strong as the syllabus is only a guideline. Double evidence might be employed

despite non-mastery of the elimination strategy. Hence, our current plans call for redefining $C(S)$ to decrease in proportion to how far a skill is from the student's current knowledge state.

$$C(S) = \frac{1}{D(S)} \quad \text{where } D(S) \text{ is the distance of } S \text{ from the farthest known skill.}$$

$D(S)$ is the distance from the farthest, not nearest known skill since the use of S depends on all the skills linked to it. S may be linked to several skills, all but one of which are known. The unknown skills then becomes the critical piece of knowledge.

For example, consider again a player moving to cave 0 in situation 1. The implicit evidence rule P-I2 is triggered by the apparent use of double evidence P2. But with respect to overlay model 1, the earlier skill L3 has not yet been learned. Hence, $D(S) = 2$ and therefore the change to USED and APPROPRIATE is reduced by 50%. The Psychologist is being cautious in interpreting this apparently advanced behavior as evidence of a non-local improvement in the player's skill. However, this possibility is not ignored.

Explicit Evidence: Another source of evidence can be obtained from the player's response to questions asked by the Tutor. This capability is not currently implemented in WUSOR-II. We have plans to implement a facility for the Tutor to obtain explicit evidence by asking the student two types of questions: test cases and follow up questions.

In a test case question, the tutor will ask the student to order the moves for the current board state or a test case. Analyzing the response reduces to the Implicit Evidence case, except that there is a larger window into the player's reasoning. The Psychologist need not guess that the student has overlooked superior moves and rejected inferior moves: the evidence is explicit in the requested ordering. The possibility that the student has forgotten to consider one alternative (which might happen in a complex game situation) is precluded. For example, situation 1 might serve as a test

case in conjunction with the following question:

Which of the following statements do you agree with most:

- (1) Caves 0, 2 and 14 are equally safe.*
- (2) Caves 0 and 2 are equally safe, but cave 14 is more dangerous.*
- (3) Cave 0 is safer than both 2 and 14.*

The second kind of explicit evidence will be derived from follow up questions that ask the student to choose among a set of possible rationales for why the current move was chosen. Rule P-E1 will monitor this source of evidence.

P-E1: If a player chooses the wrong rationale, then increment APPROPRIATE by 1 for each skill S involved in the correct rationale but absent in the chosen rationale.

For example, a follow up question to a move to cave 0 in situation 1 might be:

Which of the following explanations apply:

- (1) Caves 0, 2 and 14 are equally safe.*
- (2) Caves 0 and 2 are equally safe, but cave 14 is more dangerous because there is double evidence for pits in 14 and only single evidence for 0 and 2. Otherwise 0 and 2 are the same.*
- (3) Cave 0 is safer than 2 because there is a bat in 2 but no bat in 0.*

Background Evidence: Every teacher has expectations about the performance of a student on the basis of that student's background. This estimate changes as experience with the student is acquired, but it provides a useful starting point.

In the first implementation of the WUSOR coach, the Psychologist asked the player to classify himself his level of skill as either "novice", "amateur", "advanced" or "expert". Each of these skill levels corresponded to a different initialization for the overlay model.

We are currently experimenting with a set of background rules that associate different starting states for the overlay model with different replies to a

questionnaire presented to the player at the beginning of his first game. These rules are triggered by the player's age and experience with the game. For example, the three rules for a secondary school player are:

P-C1: If the player is in secondary school with no previous experience, then initialize the K model to AMATEUR, i.e. familiarity with the skills L1 (+ evidence) and P1 (= likelihood).

P-C2: If the player is in secondary school and has had 1-10 games experience without coaching, then initialize the K model to ADVANCED, i.e. assume familiarity with L1, L2, L3 and P1.

P-C3: If the player is in secondary school and has had over 10 games experience without coaching, then initialize the K model to EXPERT, i.e. assume familiarity with all L1, L2, L3, P1 and P2.

Similar rules are used for pre- and post-secondary school players. The rules associate naturally bounded portions of the syllabus (as determined by dependency and complexity criteria) to various age and skill backgrounds. We do not yet have enough experience with these background rules to know whether the categories of experience we have chosen are reasonable. We plan to acquire this experience studying whether the implicit and explicit rules find a particular background skill estimate, on the average, too high or too low for players of a given background.

4. The P Critic

The Psychologist maintains a history of changes to the USED and APPROPRIATE variables in order to detect inconsistencies and discontinuities. Inconsistencies are evidence that the P rules are failing to model the student properly, while discontinuities are indications of a change in the players knowledge state. The P Critic makes these decisions.

Fig. 6 is a history graph for skill S. The graph is ideal in the sense that the player consistently fails to use skill S in situations judged appropriate by the




FIG. 6 - GRAPH FOR SKILL S

Expert, until point X at which he thereafter consistently employs the skill. There is no occasional use of the skill. The P Critic would set KNOWN to T shortly after point X.

Real situations are not this clear cut; hence a certain tolerance is allowed as shown in fig. 7. A slope of zero to 10 degrees results in KNOWN being set to NIL. A slope of 35 to 45 degrees is sufficient for the critic to set KNOWN to T. Between these two regions, KNOWN is set to ?.

"?" reflects uncertainty on the part of the Psychologist. The student may be in the process of acquiring the skill, and not yet able to use it consistently. Or the P rules may be failing to model the student properly.

When KNOWN = ?, the Tutor module of the Coach becomes cautious about assuming that the student knows the skill even in situations where the student chooses the proper move. Explicit evidence is sought by means of follow up questions inquiring about the student's rationale. In the event that no clarification is obtained, i.e. the student is inconsistent even on these questions, the Tutor will ultimately ignore the Psychologist on this skill. The result is that the Coach is reduced to providing

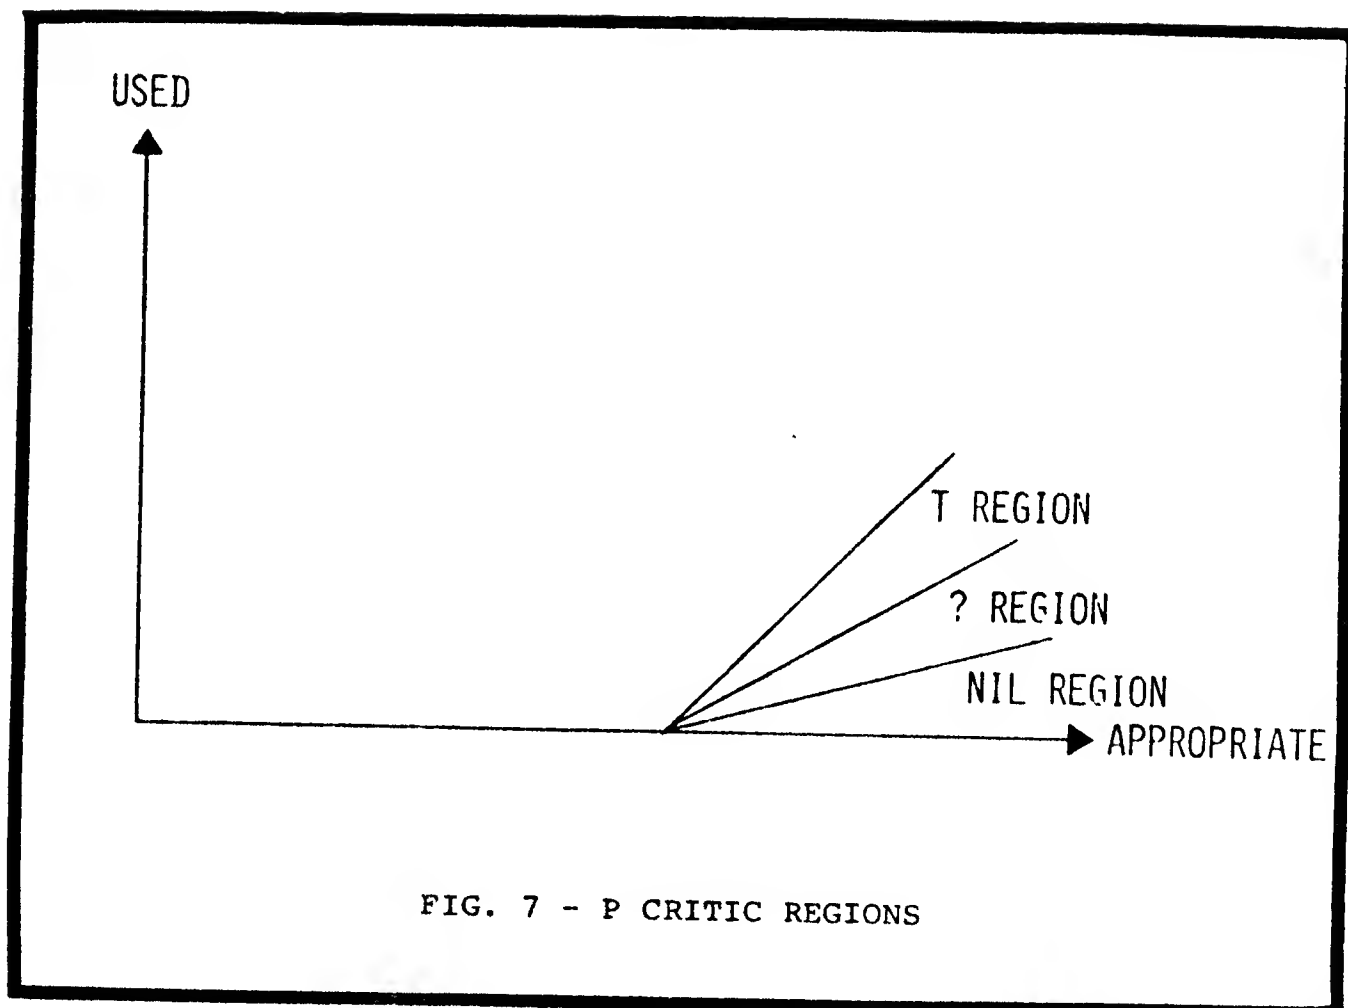


FIG. 7 - P CRITIC REGIONS

explanations generated by the Expert (when the student makes a non-optimal move) that are unpruned with respect to this skill.

5. Limitations

The modelling being conducted by the Psychologist rests on the assumption that the skills employed by the student are a subset of those of the Expert. This is not inevitable for at least three reasons. First, the student may be solving problems in a fashion completely divergent from the Expert -- there can be multiple paradigms for the particular problem domain. Second, the student may be using a non-optimal method for his own reasons. A Wumpus player may be more concerned with finishing quickly than avoiding risk, and hence choose a move to a more informative cave, despite greater risk. Third, the student may possess a skill of the Expert in an incorrect form, perhaps using it inappropriately.

We have sought to make the Expert a useful foundation for modelling by imposing certain design criteria on its design. The major one is the use of a rule system to represent the heuristics commonly employed by skilled players. This approach has been profitably employed in the medical domain [Shortliffe 74] and we similarly

find it a useful framework in which to modularly represent human skill. By interviewing skilled players and by introspection, we have evolved a rule system whose reasoning is acceptable to skilled players as capturing the essential ingredients of their own analyses. In this fashion, we have constructed an Expert for the game of Wumpus that provides a reasonable basis for modelling.

For the restricted decision making environment of Wumpus, we have not encountered multiple problem solving paradigms nor have we found it common for a student to ignore the basic strategy of choosing the safest unvisited cave. However, for other domains such as mathematical problem solving, the possibility of multiple models of expertise exists. It is a fundamental limitation of overlay modelling that a player cannot be modelled who employs a logic not understood by the Expert. Indeed, a human teacher cannot understand a student reasoning in a legitimate fashion unknown to the teacher. The power of a successful teacher arises from knowing multiple means for solving a given problem, and hence being sensitive to the particular choice made by the student. The same possibility is available to the Coach, if a Meta-Expert is provided. A Meta-Expert is a set of Experts for the given task, each modular, articulate, and comprehensible; and each capable of supplying a move analysis from its own perspective.

With a Meta-Expert, the Psychologist can attempt to identify which Expert the Student most closely approximates. The evidence distinguishing the experts derives from those situations where the predictions of the Experts differ. However, the cost of multiple experts is one more source of uncertainty. We have avoided this difficulty to date by choosing a tutoring situation -- Wumpus -- where there is broad agreement upon the part of Expert players as to the necessary skills. The design of Coaches with a Meta-Expert module is a future research goal.

Meta-experts, however, do not address the modelling difficulties arising when the student employs a skill in an incorrect form. For example, we have found some students to employ the positive and negative evidence skills for bats and pits but not for the Wumpus. The reason presumably is the greater simplicity resulting from the fact that bat and pit warnings propagate only one cave, while the Wumpus warning propagates two

caves. We have addressed this problem by not organizing the Expert around the most general set of skills. Rather positive and negative evidence has been represented by "micro" skills, one set for the 1-cave warnings of bats and pits and the other set for the two-cave warnings of the Wumpus. Our philosophy has been to break the skill analysis into sufficiently simple rules that a model which only records their presence or absence is sufficient.

It would be better to have a general theory of learning that suggested typical bugs that might occur in learning a given skill. In other research, the Coach project has studied the theory of bugs in relation to different kinds of plans [Miller & Goldstein 1976]. But in Wumpus the overall plan is simple -- find the relative dangers of the cave. Hence, we find that we are able to model the student without an elaborate bug analysis. Future research will seek to couple a theory of debugging to the theory of overlay modelling.

Given this analysis of the fundamental assumptions of overlay modelling and its limitations, there are clearly four situations where such modelling will fail. These are situations in which the underlying assumptions of these modelling rules are violated.

1. Extreme Inconsistency on the part of the player: the P critic will ultimately set the KNOWN variable of all skills to "?".
2. Unrecognized Expertise employed by the player: again the P critic will ultimately turn off the Psychologist, unless a Meta-Expert is available.
3. Player Explanations in Complex Verbal Form: natural language comprehension in the Coach is not yet implemented. Explanations expressed in English by the player are not allowed.
4. Distinguishing first order from second order bugs, that is, distinguishing the complete absence of a skill from its inappropriate use. Test question help in this situation, but are not always sufficient.

However, these situations would also task the abilities of a human teacher.

Despite these limitations, overlay modelling remains useful for two reasons. The first is that overlay modelling in its relation to explanation is essentially a linguistic theory of the Speaker. Each of us, when formulating an explanation, abbreviates the explanation in accord with our model of the listener. This model is based on our analysis of the listener's behavior in terms of the knowledge we believe is relevant. Overlay modelling performs a similar function for the Coach. A human speaker or computer coach may have a mistaken model of the listener, but ultimately a person or computer can judge another only in terms of what he, she or it knows itself.

The second reason arises from the special demands of the educational context. The Coach is not an impartial observer, but rather has the goal of conveying its style of expertise. Hence, its insight into the student can be useful, even if limited to hypotheses regarding which aspects of its expertise the student possesses. Its goal is to convey that style it knows about; its modelling is to determine how much of that style is known.

6. Experimental Program

The fundamental question is how accurate are the K and L models as estimates of the player's knowledge and learning preferences. To address this question, we are employing 4 different classes of experiment.

1. Turing Tests: Human players will be analyzed by interviewers to provide benchmarks for the level of modelling that can be achieved by competent human teachers. In one variation, an *accomplice* will be asked to deliberately simulate certain student strategies, and the ability of human observers to detect these strategies will be studied. These Turing Tests will determine if the Psychologist module provides modelling performance comparable to human observers.
2. Articulate Psychologist Experiments: Our rule-based approach to modelling allows the Psychologist to explain its hypotheses by reporting which rules were triggered and by what evidence. The accuracy of these self-explanations will be judged by both the student himself, and an interviewer who observes the student's play and discusses his moves with him.

3. Closed Loop Experiments: The game will be played by a modified version of the Expert program which employs a sub-optimal strategy. The Psychologist will be judged by whether it diagnoses the strategy.
4. Predictive Experiments: An overlay model can yield a deterministic procedural model of a player by deleting all rules of the Expert with KNOWN = NIL. The result is a simulated player that can be used to predict the player's performance. The accuracy of these predictions will provide another test of the psychologist's success.

To date, we have carried out informal "articulate psychologist" experiments with Wusor-II. Players over a wide spectrum of skill find the comments generated by the psychologist module to be comprehensible and reasonable, as evaluated by interviews with the players. We have also run closed loop experiments in which an impartial player consistently employs a sub-optimal strategy. WUSOR-II successfully diagnoses this. We are currently in the process of designing simulated players to serve as rigorous closed loop tests.

We plan over the next 12 month period to run the two most ambitious classes of experiments, Turing Tests and Predictive Experiments. Our subject population will be undergraduates enrolled in an education major. (We will be interested both in the success WUSOR has in coaching these students and in their reactions to WUSOR as an educational tool.)

In summary, we are encouraged by reactions of students and teachers to the current state of the Coach, but rigorous evaluation of overlay modelling remains to be done.

7. Related Literature

WEST: The WEST program by Burton and Brown [1976] is a computer coach for the PLATO game "HOW THE WEST WAS WON". In this game, a player must form from three numbers an arithmetic expression whose value is either the largest possible, or occasionally a given number. The educational purpose of the game is to provide experience with arithmetic operators and the use of parentheses. C. Resnick [1975] found that many students reached plateaus, such that they failed to improve their skill, although they

continued to enjoy the game. The WEST coach was designed to discuss less than optimal moves with the student in order to move him or her off such a plateau.

Burton and Brown model the student by contrasting his or her move to the move recommended by the expert. USED and APPROPRIATE variables are maintained to record the frequency with which different skills are employed. Burton and Brown's development of this modelling technique was our starting point. We have extended their approach in three ways.

(1) A syllabus is introduced that organizes the skills in a complexity/dependency graph. For complex situations this is required. For simpler domains with a limited number of skills, it is less important. For WEST, the syllabus of fig. 8 might have been employed, which reflects the usual order in which arithmetic skills are taught.

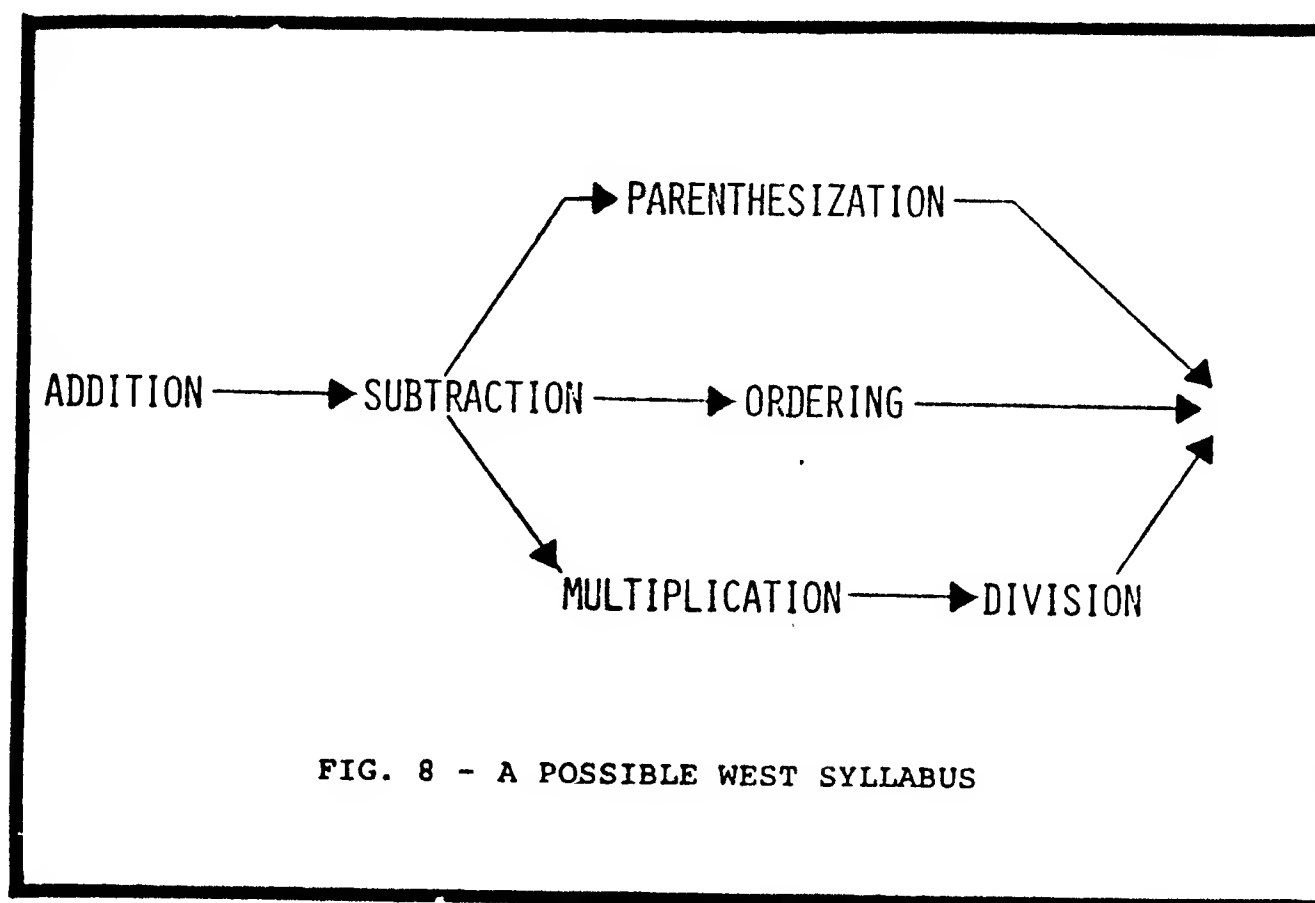


FIG. 8 - A POSSIBLE WEST SYLLABUS

(2) A P critic is introduced to observe discontinuities and inconsistencies. This is important to observe when the modelling is failing to capture the student's behavior. It could readily be applied to the WEST case.

(3) Multiple sources of evidence are used to increase the window into the student's reasoning. WEST relied solely on implicit evidence derived from the student's play. A facility for obtaining explicit evidence through follow up questions could be incorporated into the WEST coach.

BIP: BIP [Wescourt 1976] is a CAI program for tutoring elementary programming skills. We mention it here to cite an alternative to Expert-based overlay modelling. BIP uses a very detailed syllabus as does overlay modelling. But BIP associates with each skill in the syllabus (called the Curriculum Information Network) a set of specific exercises and a description of the various correct and incorrect solutions. A skill is attributed to the student if he or she succeeds at these exercises.

The virtue of this approach is that the diagnosis of whether a skill is employed is much simpler to make. An elaborate domain expert is not needed. The disadvantage, however, is that the tasks are very restrictive, e.g. a typical one might be to "PRINT A LITERAL". The greater complexity of overlay modelling with respect to an embedded Expert program is required to allow free choice by the student in more complex problem settings.

Human Problem Solving: Overlay modelling is a potentially valuable tool for information processing psychology. Hence we compare it here done by Newell and Simon [1972] and their colleagues. Overlay modelling can be used to induce a production system model of a human problem solver. The required ingredient is an Expert that analyzes the problem solver's acts in terms of a set of constituent skills -- in this case a set of productions. This notion of comparing a problem solving protocol to the behavior of a production system is briefly described as the "trace" feature of the PAS-II protocol analysis program [Waterman and Newell 1973].

In the computer coach context, we have not attempted the level of detail in modelling that Newell and Simon seek, wherein even eye movements must be accounted for. The Coach does not have that much information regarding the student's behavior. Indeed, we do not allow unrestricted English interaction. In the PAS-II protocol analysis program, English is permitted by making the program interactive -- i.e. a human analyst can aid in interpreting the protocol. Such a solution is not applicable to the real time demands made by the computer coaching context.

Our development of overlay modelling suggests an extension to production based modelling, in the form of the Syllabus. The productions for a given problem domain can

be organized into a network reflecting complexity and dependency. This network then suggests the order in which the productions are acquired.

8. Conclusions

Overlay modelling constitutes a set of techniques for describing a person's problem solving skills in terms of an expert program for the task. These techniques are rule systems for monitoring multiple sources of evidence, overlays for structuring the model, and a critic for detecting non-linearities. This approach has limitations, but it has already shown itself to be useful for maintaining a model of the learner's state as part of an AI-based CAI program.

Ultimately, progress towards an improved theory of modelling will have an important impact on the following areas:

1. In CAI by addressing the critical need to model the learner so as to provide high quality personalized instruction.
2. In education by offering overlays as a structural, non-numerical model of the student.
3. In applied AI by improving the ability of an AI program employed as an intelligent assistant to generate appropriate explanations for the user.
4. In theoretical AI by defining criteria such as comprehensibility and modularity that expert programs should satisfy if they are to be useful as part of an AI-based CAI systems.
5. In information processing psychology by developing a procedural theory for inducing models of a subject's problem solving behavior.

9. References

- Brown, J.S., R. Burton, & A. Bell, "SOPHIE: A Step Toward Creating a Reactive Learning Environment", International Journal of Man-Machine Studies, Vol. 7, 1975, pp. 675-96.
- Brown, J.S., R. Burton, M. Miller, J. DeKleer, S. Purcell, C. Hausmann, R. Bobrow, Steps Towards a Theoretical Foundation for Complex, Knowledge-based CAI, ICAI Report 2, Bolt Beranek and Newman, Cambridge, Ma, Aug., 1976.
- Burton, R. & J. S. Brown, "A Tutoring and Student Modelling Paradigm for Gaming Environments", SIGCSE Bulletin, Feb. 1976, pp. 236-246.

- Carr, B. Wumpus Advisor II, MIT AI Laboratory, forthcoming memo, Mar. 1977.
- Collins, A. & M. Grignetti, Intelligent CAI, BBN Report No. 3181, Bolt, Beranek & Newman, Inc, Cambridge, MA, 1975.
- Goldstein, I. P. The Computer as Coach: an Athletic Paradigm for Intellectual Education, MIT Artificial Intelligence Laboratory, Memo 389, February 1977.
- Miller, M. and I. P. Goldstein, Overview of a Linguistic Theory of Design, MIT Artificial Intelligence Laboratory Memo 383, December, 1976.
- Newell, A. and H. Simon, Human Problem Solving, NJ:Prentice Hall, 1972.
- Resnick, C., Doctoral Dissertation, University of Illinois, 1975.
- Shortliffe, T. 1974. MYCIN -- A Rule-based computer program for advising physicians regarding antimicrobial therapy selection. SAIL Memo 251.
- Stansfield, J., B. Carr, & I. P. Goldstein, Wumpus Advisor I: A First Implementation of a Program that Tutors Logical and Probabilistic Reasoning Skills, MIT Artificial Intelligence Laboratory Memo No. 381, Sept. 1976.
- Waterman, D.A., & A. Newell, "PAS-II: An Interactive Task-Free Version of An Automatic Protocol Analysis System," *IJCAI III*, 1973, pp. 431-445.
- Yob, G., "Hunt the Wumpus", Creative Computing, Sep-Oct, 1975, pp. 51-54.
- Wescourt, K., BIP's Curriculum Information Network, unpublished report, Oct. 1976.

